

# Squarepants in a Tree: Sum of Subtree Clustering and Hyperbolic Pants Decomposition

David Eppstein\*

## Abstract

We provide efficient constant factor approximation algorithms for the problems of finding a hierarchical clustering of a point set in any metric space, minimizing the sum of minimum spanning tree lengths within each cluster, and in the hyperbolic or Euclidean planes, minimizing the sum of cluster perimeters. Our algorithms for the hyperbolic and Euclidean planes can also be used to provide a *pants decomposition*, that is, a set of disjoint simple closed curves partitioning the plane minus the input points into subsets with exactly three boundary components, with approximately minimum total length. In the Euclidean case, these curves are squares; in the hyperbolic case, they combine our Euclidean square pants decomposition with our tree clustering method for general metric spaces.

## 1 Introduction

A *hierarchical clustering* of a finite set of points can be visualized as a binary tree, having the points at its leaves. In such a tree, we can form a cluster for each internal node, of the points descending from it. These clusters, together with the empty set, will form a family  $\mathcal{F}$  of subsets of the points, with the property that any two subsets in the family are either disjoint or related by containment; this family is maximal in the sense that no additional set can be added to it while preserving this property. Equivalently, a family of maximal sets of this type can be viewed as forming a binary tree, with an internal node per nonempty set; each node has as its children the maximal subsets in the family. Figure 1) shows these two equivalent views of a hierarchy.

There has been much work on heuristics for hierarchical clustering of points in metric spaces [20], often based on *agglomerative clustering* methods that start from singleton clusters and repeatedly merge pairs of clusters into single larger clusters until only one set remains [10, 15, 30]. For instance, the *single linkage* clustering method is essentially equivalent to Kruskal's algorithm for minimum spanning tree construction, and the *neighbor-joining* method [26], which can be defined as a general clustering technique in this way, is widely used for reconstructing evolutionary trees.

However, there has been less work on problems of finding a clustering that optimizes some objective function measuring the overall quality of the clustering [12]. In this paper we consider problems of finding a clustering minimizing the sum of cluster sizes, where we may measure the size of a cluster either by the length of its minimum spanning tree (for general metric spaces) or by the perimeter of its convex hull (for the Euclidean and hyperbolic planes). We are unaware of prior work on these versions of the optimal hierarchical clustering problem.

We also consider related problems of cutting the plane by a system of disjoint simple closed curves such that each component of the plane minus the curves and the input points has three boundary components, minimizing the total curve length. Such a *shortest pants decomposition* has been approximated in the Euclidean case by Poon and Thite [24], and related decompositions of surfaces into simple components by short curves have proven useful as building blocks for other topological

---

\*Computer Science Department, Donald Bren School of Information & Computer Sciences, University of California, Irvine; eppstein@uci.edu




Figure 1: A hierarchical clustering of five points (left) and the corresponding binary tree (right).

computations [17]. For the Euclidean case, we provide a simple quadtree-based approximation algorithm that is less accurate but more efficient than that of Poon and Thite. We also provide similar approximation algorithms for pants decomposition in the hyperbolic plane; to our knowledge this version of the pants decomposition problem has not been studied previously.

This paper is independent of, and concurrent with, a recent paper by Krauthgamer and Lee [22], which claims to be the first to study approximation algorithms in hyperbolic spaces. Krauthgamer and Lee obtain a polynomial time approximation scheme for the traveling salesman problem in any fixed dimensional hyperbolic space, by a technique very similar to our method for hyperbolic clustering, in which points in low-diameter clusters are approximated with Euclidean spaces while the connections between the clusters are approximated by trees.

## 2 New Results

We prove the following results.

- We formulate the problem of hierarchical clustering minimizing the sum of spanning tree lengths of the clusters, for general metric spaces, show that it is NP-complete, and provide a constant factor approximation algorithm for the problem.
- We formulate the problem of hierarchical clustering in the Euclidean plane, minimizing the sum of convex hull perimeters, and relate it to the previously studied problem of optimal pants decomposition [24]. We provide a simple example showing that the two problems do not always have equal solutions, but we show that they can both be approximated to within a constant factor of the optimal total length by a simple quadtree-based clustering algorithm, in time  $O(n \log n)$ .
- By analogy with the Euclidean case, we formulate the sum-of-perimeter clustering and pants decomposition problems in the hyperbolic plane. We provide an approximation algorithm based on a combination of our Euclidean technique (for point sets with diameter  $O(1)$ ) and our general metric space technique (for point sets with closest distance  $\Omega(1)$ ). Our hyperbolic approximation uses a lemma that may be of independent interest: for any hyperbolic point set with closest distance  $\Omega(1)$ , the convex hull and minimum spanning tree have lengths within a constant factor of each other.

## 3 Hardness of Sum of Subtree Clustering

Define an  $i$ -bisectable tree, for integer  $i \geq 0$ , as follows: a 0-bisectable tree is just a tree with a single vertex, and an  $i$ -bisectable tree, for integer  $i > 0$ , is formed by connecting any two  $(i - 1)$ -bisectable trees by a single edge connecting any two of their vertices. An  $i$ -bisectable tree always has exactly  $2^i$




Figure 2: The three 3-bisectable trees.

vertices. We say that a tree is *bisectable* if it is an  $i$ -bisectable tree for some  $i$ . Up to isomorphism of free trees, there is only one 1-bisectable tree (a single edge), one 2-bisectable tree (a path of four vertices), and three 3-bisectable trees (Figure 2).

Any  $i$ -bisectable tree either has at most a single nontrivial automorphism that exchanges its two  $(i-1)$ -bisectable subtrees. Based on this observation, if we let  $d_i$  represent the number of  $i$ -bisectable trees,  $s_i$  denote the number of symmetric  $i$ -bisectable trees, and  $a_i = d_i - s_i$  the number of asymmetric  $i$ -bisectable trees (e.g.,  $d_3 = 3$ ,  $s_3 = 2$ , and  $a_3 = 1$ ), we can compute these numbers by the following recurrence.

$$\begin{aligned} s_i &= 2^{i-1}a_{i-1} + 2^{i-2}s_{i-1} \\ a_i &= \binom{s_i}{2} \\ d_i &= a_i + s_i \end{aligned}$$

Using this recurrence we counted 136 4-bisectable trees, and 2098176 5-bisectable trees.

LEMMA 3.1. *We can test whether an  $n$ -node tree is bisectable in time  $O(n)$ .*

*Proof.* Given a tree  $T$ , perform the following steps:

- Choose a root for  $T$  arbitrarily.
- For each vertex  $v$  of  $T$ , count the number of descendants of  $v$  (including  $v$  itself). This can be done by a single postorder traversal of  $T$ , as the number of descendants of  $v$  is one plus the sum of the numbers of descendants of each of its children.
- Define an edge of  $T$  to be *odd* if the farther of its two endpoints from the tree root has an odd number of descendants. Identify the set of odd edges.
- If  $T$  does not have exactly  $n/2$  odd edges then  $T$  is not bisectable. Otherwise, form a tree  $T'$  with  $n/2$  vertices by contracting each odd edge of  $T$ .  $T$  is bisectable if and only if  $T'$  is bisectable, which can be tested by a recursive application of the same algorithm.

If  $T$  is bisectable, consider the sequence of forests of  $2^i$  subtrees formed by splitting  $T$  according to the first  $i$  levels of the bisection hierarchy of  $T$ . Then it is straightforward to show by induction on  $i$  that, at each step of this sequence before the last, the set of edges that are odd in their subtrees remains unchanged. Therefore, if  $T$  is bisectable, the odd edges connect pairs of vertices in the forest of  $n/2$  subtrees at the bottom level of the hierarchy, and contracting these edges does not change the bisectability of the remaining hierarchy; thus, the algorithm will report correctly that  $T$  is bisectable.




Figure 3: Reduction for NP-completeness of finding bisectable subtrees. On the upper left is a twelve-vertex graph  $G$  for which we wish to solve the  $H$ -matching problem; lower left is a path used to pad its size to a power of two. On the right we have the paths  $p_i$ , with their endpoints  $u_i$  connected to each other. The edges between  $u_i$  and the vertices on the left are shown light and dashed in order to avoid obscuring the drawing.

Conversely if the algorithm reports that  $T$  is bisectable, a bisection hierarchy for  $T$  can be constructed from the corresponding hierarchy for  $T'$ . Thus, the algorithm correctly reports the bisectability of  $T$ .

Each step of the algorithm except for the recursive call takes linear time. The sizes of the subtrees passed as arguments to recursive calls shrink in a geometric series, so the total time for the overall algorithm is also linear.  $\square$

**THEOREM 3.1.** *It is NP-complete, given an undirected graph  $G$  with  $2^i$  vertices, to determine whether  $G$  has an  $i$ -bisectable subtree.*

*Proof.* Membership in  $NP$  follows since we can use Lemma 3.1 to test whether a given subtree is bisectable. To show that the problem is  $NP$ -complete, we reduce from the known  $NP$ -complete problem of  $H$ -matching (that is, covering all vertices of a graph by disjoint copies of a fixed subgraph  $H$ ).  $H$ -matching is known to be  $NP$ -hard for all connected subgraphs  $H$  with more than two vertices [19, 21]; in our case, we use as  $H$  the unique 3-bisectable tree; that is, the path on four vertices.

So, suppose we are given a graph  $G$ , and wish to determine whether  $G$  has an  $H$ -matching. We reduce the problem to the existence of a bisectable subtree on a larger graph. We can assume without loss of generality that the number of vertices in  $G$  is bisectable by four; otherwise  $G$  can have no  $H$ -matching. Let  $G'$  be the disjoint union of  $G$  with paths of length four, sufficient so that the number  $n$  of vertices in  $G'$  is a power of two. Finally, form graph  $G''$  by adding to  $G'$   $n$  additional vertices, in the form of  $n/4$  paths  $p_i$ . Choose an endpoint  $u_i$  of each such path, and add additional edges between every pair of vertices  $u_i, u_j$  and between each  $u_i$  and each vertex of  $G$ . The completed graph  $G''$  is depicted in Figure 3.

We claim that  $G$  has an  $H$ -matching if and only if  $G''$  has an  $i$ -bisectable subtree. For, if  $G$  has an  $H$ -matching, then we can cover  $G''$  by paths of length eight by connecting each  $p_i$  to one of the paths in the  $H$ -matching; we can then merge these paths into a  $i$ -bisectable tree using the edges between

pairs  $u_i, u_j$ . Conversely, if  $G''$  has an  $i$ -bisectable subtree, then by repeatedly partitioning that subtree according to the definition of divisibility, we find a cover of  $G''$  by paths of length four. Each path  $p_i$  must be in that cover (because it is the only path that covers the endpoint most distant from  $u_i$ ) so the remaining paths must form an  $H$ -matching of  $G$ .

Thus we have reduced  $H$ -matching to our  $i$ -bisectable subtree problem, completing the proof of  $NP$ -completeness of that problem.  $\square$

The same reduction shows more generally that, given  $G$  and  $i \geq 2$ , it is  $NP$ -complete to determine whether  $G$  can be covered by disjoint  $i$ -bisectable subtrees.

**THEOREM 3.2.** *It is  $NP$ -complete, given a metric space  $\mathcal{M}$  (specified as its distance matrix) and a number  $K$ , to determine whether  $\mathcal{M}$  has a hierarchical clustering in which the sum of minimum spanning tree lengths of the clusters is at most  $K$ .*

*Proof.* Membership in  $NP$  is straightforward, since we may demonstrate that there exists a clustering with small total size by exhibiting the clustering.

To prove  $NP$ -hardness, we reduce from the problem of finding an  $i$ -bisectable subtree, which we have seen is  $NP$ -complete. Suppose we are given a graph  $G$ , with  $2^i$  vertices, and wish to determine whether  $G$  contains an  $i$ -bisectable subtree. From  $G$  we form a metric space  $\mathcal{M}$  in which the points are the vertices of  $G$ ; two points are at distance one if they are connected by an edge in  $G$ , and at distance two otherwise. We set  $K = i2^i - 2^i + 1$ .

If  $G$  contains an  $i$ -bisectable subtree  $T$ , we form a clustering by recursively splitting  $T$  into smaller bisectable subtrees. Each cluster in this clustering has a spanning tree with unit length edges, namely the subtree of  $T$  induced by its vertices, so we can calculate that the total length of all minimum spanning trees of clusters is exactly  $K$ . Conversely, any clustering of length  $K$  must have  $2^{i-j}$  vertices in each cluster at level  $j$  of the clustering, or else even if all spanning tree edges have unit length the total length would exceed  $K$ . Further, each pair of clusters at level  $j$  must be connected by at least one edge of  $G$ , or else some spanning tree would have an edge of length two, again causing the total length to exceed  $K$ . So, in this case, choosing an edge connecting each pair of clusters at each level of the tree gives us an  $i$ -bisectable subtree of  $G$ .

Thus we have reduced the  $i$ -bisectable subtree problem to our clustering problem, completing the proof of  $NP$ -completeness of that problem.  $\square$

#### 4 Approximate Sum of Subtree Clustering

Suppose that we wish to find a hierarchical clustering of a set of points in a metric space, approximately minimizing the sum of cluster sizes, where we measure the size of a cluster by the total length of the edges in its minimum spanning tree.

We observe the optimal clustering may have a structure quite different from that of the minimum spanning tree of  $\mathcal{M}$ . For one thing, our  $NP$ -completeness reduction shows that, if one forms a graph connecting points at close distances to each other, it may be more important to find a bisectable spanning tree in this graph than a graph of minimum possible weight. For another, even in a tree metric, if we restrict ourselves to partitions formed by recursively splitting the tree on its edges we may form a clustering very far from optimal; for instance, for the star  $K_{1,n-1}$  (with unit edge lengths) recursive splitting produces a clustering with total weight  $\Omega(n^2)$  while the optimal clustering has total weight  $O(n \log n)$ .

Nevertheless, as we now show, the idea of recursively splitting the minimum spanning tree can lead to an efficient approximation for the optimal clustering. We approximately cluster  $\mathcal{M}$  by the following algorithm, the operation of which is depicted in Figure 4.




Figure 4: Approximation algorithm for subtree clustering. Top left: a minimum spanning tree  $T$  of a point set, with edges labeled by length. Top right: tree  $T^*$  subdivided from  $T$  so that each vertex has degree at most three; new vertices are unlabeled. Center: partitioning  $T^*$  to minimize the maximum length of the two resulting subtrees. Bottom: the clustering formed by continuing the partition recursively in each subtree.

1. Compute a minimum spanning tree  $T$  of the points.
2. While  $T$  contains a vertex  $v$  with degree greater than three, split  $v$  into two vertices connected by an edge of length zero, each adjacent to two or more of the neighbors of  $v$ . Give one of these two vertices the identity of the original input point it came from.
3. Call the resulting tree resulting from these split operations  $T^*$ ; it has at most three edges per vertex.
4. Find the edge  $e$  of  $T^*$  such that the maximum total length among the two subtrees  $T_\ell^*$  and  $T_r^*$  is as small as possible.
5. Form two clusters consisting of the input points corresponding to vertices in these two subtrees.
6. Use the two subtrees  $T_\ell^*$  and  $T_r^*$  to partition these two clusters recursively.

To prove that this algorithm produces an approximation to the optimal clustering, we prove an upper bound on the total weight of the clustering produced by the algorithm, and a matching lower bound on any clustering. The upper bound depends on a standard lemma on tree separators.

**LEMMA 4.1.** *In any tree with nonnegative edge weights, let  $e$  be an edge the removal of which minimizes the maximum weight among the two resulting subtrees. Then the weight of each subtree formed by the removal of  $e$  is at most  $2/3$  of the total weight of the initial tree.*

*Proof.* Let the tree be  $T$  and its weight be  $W$ . If  $e$  is an edge separating  $T$  into subtrees one of which has weight larger than  $(2/3)W$ , let  $f$  be the edge adjacent to  $e$  in this heavy subtree, such that the weight of the subtree containing  $f$  is larger than the weight of the other subtree adjacent to the same vertex of  $e$ . Then the subtree containing  $f$  has weight at least  $W/3$ , so  $f$  forms a better split: it partitions the tree into two subtrees, one of which (the one containing  $e$ ) has weight at most  $(2/3)W$ , and the other of which either has less weight or fewer edges than the heavy subtree for  $e$ . We can not find an infinite sequence of better splits, so there must be an edge forming a good split as described by the lemma.  $\square$

**LEMMA 4.2.** *Let the edges of the minimum spanning tree  $T$  have weights  $w_0, w_1, \dots$  in descending order, and let  $W = \sum w_i$ . Then the total weight of the spanning trees of the clusters produced by our clustering algorithm is at most  $\sum w_i(1 + \log_{3/2}(W/w_i))$ .*

*Proof.* An edge  $e_i$  of  $T$  with weight  $w_i$  participates in the minimum spanning trees of clusters at levels  $0, 1, 2, \dots$  of the clustering, until reaching a level in which its two endpoints are in different clusters. At level  $k$ , the total weight of the spanning tree in which it participates is (by Lemma 4.1) at most  $W(2/3)^k$ ; if  $k > \log_{3/2}(W/w_i)$ , this total weight is less than  $w_i$ , so  $e_i$  can only participate in trees in at most  $\log_{3/2}(W/w_i)$  levels of the hierarchy.  $\square$

**LEMMA 4.3.** *For any decreasing sequence of values  $w_i$  with sum  $W$ ,*

$$\sum w_i \left( \frac{1}{2} + \lfloor \log_2(i+1) \rfloor \right) \geq \sum \frac{1}{2} w_i \log_2(W/w_i).$$

*Proof.* By scaling the  $w_i$  by the same factor, we may assume without loss of generality that  $W = 1$ . We interpret the  $w_i$  as probabilities of drawing symbol  $i$  in a random variable that takes the indices  $0, 1, \dots, n-1$  as its values, with index  $i$  having probability  $w_i$ . There exists a binary code (that is, a binary tree having these indices at its leaves) in which the expected path length to a randomly drawn index is  $1 + \sum 2w_i \lfloor \log_2(i+1) \rfloor$  (Figure 5). The result then follows from Shannon's entropy lower bound of  $\sum w_i \log_2(1/w_i)$  on the average path length of any binary code [27].  $\square$




Figure 5: A binary code with expected path length  $1 + \sum 2w_i \lfloor \log_2(i+1) \rfloor$ .

The relation between this sum of logs of ranks and entropy is closely related to the efficiency of one-to-one codes in coding theory [9, 23]; codes similar to the one shown in Figure 5 have also been used as part of more complex data compression schemes [6, 13]. We note that an inequality in the other direction,  $\sum w_i \lfloor \log_2(i+1) \rfloor \leq \sum w_i \log_2(W/w_i)$ , follows more trivially:  $i+1 \leq W/w_i$ , since otherwise the sum of the first  $i+1$  weights would exceed  $W$  [29].

LEMMA 4.4. *Let  $C$  be any hierarchical clustering in a metric space  $\mathcal{M}$ , let the edges of the minimum spanning tree of  $\mathcal{M}$  have weights  $w_0, w_1, \dots$  in descending order, and let  $W = \sum w_i$ . Then the total weight of the minimum spanning trees of the clusters in  $C$  is at least  $\sum \frac{1}{2} w_i (1 + \log_2(W/w_i))$ .*

*Proof.* Let  $C_i$  be the total length of the minimum spanning trees of the clusters at level  $i$  of the clustering, and let  $F_i$  be the total length of the minimum forest in  $\mathcal{M}$  having  $2^i$  trees. Note that this minimum forest can be computed by removing the  $2^i - 1$  largest edges from the minimum spanning tree of  $\mathcal{M}$ . Then

$$\begin{aligned} \sum C_i &\geq \sum F_i \\ &= \sum_i \sum_{j \geq 2^i - 1} w_j \\ &= \sum w_i (1 + \lfloor \log_2(i+1) \rfloor) \\ &\geq \sum \frac{1}{2} w_i (1 + \log_2(W/w_i)), \end{aligned}$$

where the final inequality in this sequence is Lemma 4.3. □

THEOREM 4.1. *If we are given a minimum spanning tree of a metric space  $\mathcal{M}$ , we may compute in  $O(n \log n)$  time a hierarchical clustering, such that the sum of minimum spanning tree lengths of clusters in our clustering is within a factor of  $2 \log_{3/2} 2 \sim 3.42$  of the optimal clustering.*

*Proof.* We may implement the approximation algorithm described above by using the dynamic tree median data structure of Alstrup et al. [3] to find the edge to be removed at each step of the recursive partition of the tree  $T^*$ . This data structure takes time  $O(\log n)$  per step, and all other operations of the algorithm may easily be implemented in total time  $O(n)$ , so the total time for the algorithm is  $O(n \log n)$ . The approximation ratio for our algorithm follows from Lemmas 4.2 and 4.4. □






Figure 6: A set of 16 points such that the clustering minimizing the total perimeter of the hulls of the clusters has non-disjoint hulls of clusters. The two top-level clusters in the optimal clustering are shown.

## 5 Euclidean Sites

In this section we consider two problems. The first is a clustering problem very similar to the one we have already studied for general metric spaces: hierarchically clustering point sites in the Euclidean plane  $\mathbb{R}^2$  in such a way as to minimize the sum of cluster perimeters. The second problem we consider is one of optimal *pants decomposition*. A *pair of pants* is a topological surface in the form of a disk with two holes cut into it; that is, it can be represented as a connected subset of the plane the boundary of which has three connected components. A *pants decomposition* is a partition of a topological surface into pairs of pants. The Euclidean pants decomposition problem considers as input a set of point sites in the plane, and asks for a family of disjoint closed curves such that removing the curves and the sites from the plane decomposes it into connected components all of which are pairs of pants. Each pair of pants must have one curve as its outer boundary, except for one infinite pair of pants having a boundary at infinity. However the two inner boundaries of each pair of pants may either be other curves or input sites. Figure 1 (left) and Figure 4 (bottom) depict points in the plane of the drawing, surrounded by curves; if one removes the outer curve from each of these figures, they can be interpreted as pants decompositions.

The question of minimum-length pants decomposition was proposed by Jeff Erickson and Kim Whittlesey (personal communication) and first attacked algorithmically by Poon and Thite [24], who provided both a polynomial time approximation scheme and a polynomial time exact algorithm for a restricted version of the problem in which the curves must be rectangles.

### 5.1 Clustering versus Pants

Clustering and pants decomposition are closely related. We can form a clustering by creating a cluster for the points within each curve of a pants decomposition, and these clusters (together with singleton clusters for each site and a cluster for all sites) form a hierarchical clustering. The length of the curve surrounding each clustering is at least the perimeter of the convex hull of the cluster. Therefore, the length of a pants decomposition, plus the length of the convex hull of all the sites, is lower bounded by the length of the clustering that minimizes the sum of perimeters. However, the two problems may have total lengths that differ by arbitrarily large factors, due to the inclusion of the whole set in the clustering and not in the pants decomposition; for instance, the three points  $(0,0)$ ,  $(0,1)$ , and  $(L,0)$  (for large  $L$ ) have a pants decomposition with length  $2 + \epsilon$  for any  $\epsilon$  (formed by a curve around the two closest point) while its optimal clustering has length  $2L(1 + o(1))$ .

Even ignoring the presence or absence of a cluster for the entire point set, the two problems can




Figure 7: A quadtree and the corresponding compressed quadtree.

differ in other ways. Figure 6 depicts a set of 16 points for which the unique optimal clustering has two clusters with overlapping convex hulls. We computed the optimal clustering using a dynamic programming algorithm that determines the optimal clustering for each set of points, by considering all partitions of that set into two previously clustered subsets, in total time  $O(3^n)$ .

We do not know whether the optimal pants decomposition of these points has the same hierarchical structure, but it must differ in total length. We can also show that there exist point sets for which the optimal clustering and the optimal pants decomposition differ in hierarchy: either the depicted point set is such an example, or we can find such an example by scaling the vertical coordinates of this point set by a factor  $0 < s < 1$  while leaving the horizontal coordinates unchanged. As we scale the points, for sufficiently small  $s$  the points become sufficiently close to colinear that their optimal clustering has disjoint hulls; let  $s_0$  be the largest  $s$  for which this is true. Then scaling by  $s_0$  produces a point set in which two clusterings, one with disjoint hulls and one with non-disjoint hulls, have equal lengths, however the pants decomposition corresponding to the clustering with disjoint hulls is strictly shorter than the pants decomposition corresponding to the clustering with non-disjoint hulls. Scaling by  $s_0 + \epsilon$  for small epsilon can therefore be seen to have a different hierarchy in its optimal clustering and its optimal pants decomposition. The non-disjoint hulls in this example also imply that it may be difficult to find a polynomial time approximation scheme for the optimal clustering problem, as the technique used by Poon and Thite to approximate the optimal pants decomposition relies on the disjointness of the curves it finds.

## 5.2 Euclidean Clustering

In this section we describe a fast and simple quadtree-based heuristic for finding an approximation to the Euclidean clustering problem. We will later show how to adapt the same algorithm to find similar approximations for pants decomposition. Our algorithm will turn out to be an important subroutine in our later approximation algorithm for similar problems in the hyperbolic plane.

Our fast Euclidean approximation technique is based on the *compressed quadtree* [4, 7, 8, 11, 16]. A quadtree is a well-known recursive space partition data structure formed from a set of sites by surrounding the sites by a square bounding box, then, as long as some minimal square of the structure contains more than one input site, splitting those squares into four smaller squares. To handle cases

of ambiguity, we consider a point on the boundary of a square to be outside the square if it is on the lower or right sides of the square, and inside the square otherwise; in this way each square is exactly the disjoint union of its four quadrants. A compressed quadtree is a subset of the squares of a quadtree, consisting only of those squares that contain sites in more than one of their four quadrants. Figure 7 depicts a quadtree and compressed quadtree for a set of sites. Each square of a compressed quadtree other than the initial root square is contained in a unique minimal larger square of the compressed quadtree, which we consider to be its parent. We also consider each site to be a node in the structure, with parent the minimal square containing it. Thus, the structure forms a tree, in which the leaves are the input sites, and each internal node corresponds to a square with at least two children, at most one child per quadrant of the square. Since each square has two or more children, there are at most  $n - 1$  squares in the compressed quadtree. A compressed quadtree for a set of  $n$  sites may be built or maintained dynamically in  $O(n \log n)$  time, in a model of computation in which we may perform arithmetic and bitwise Boolean operations on the binary representations of the site coordinates [8, 16].

To form a hierarchical clustering from a compressed quadtree, we form a cluster for the set of sites inside each square of the quadtree. This is not itself a hierarchical clustering, because some squares in the quadtree may have more than two children. If a square has three children, we also form a cluster combining the sites in two of its children, choosing two out of the three children whose quadrants share a side with each other. If a square has four children, we form two clusters combining the children in adjacent pairs. As the set of sites inside a square has a convex hull perimeter bounded above by the perimeter of the square, the total convex hull perimeters of the clusters in the resulting hierarchical clustering is bounded above by the total perimeters of the squares in the compressed quadtree. Note that, in this clustering, disjoint clusters have disjoint convex hulls.

We analyze this clustering using an integral involving *local feature size* [5, 25], reminiscent of Rupert's similar analysis of the optimal number of triangles in a bounded-aspect-ratio triangulation [25]. Given a set  $S$  of input sites, define the local feature size  $\text{lfs}(x, y)$  to be the distance from point  $(x, y)$  to the second nearest point in  $S$ . Let  $\square$  denote the root square of our quadtree, which we assume to be a minimal bounding square for the input points; thus, for points  $(x, y) \in \square$ ,  $\text{lfs}(x, y)$  is bounded above by the diagonal length of the square and below by half the minimal separation between any two points.

LEMMA 5.1. *The total perimeter of the clusters in our clustering is  $O(\int_{(x,y) \in \square} 1/\text{lfs}(x, y) dx dy)$ .*

*Proof.* We prove more generally that, in the (non-compressed) quadtree for  $S$ , the sum of perimeters of the squares  $C_i$  is  $O(\int_{(x,y) \in \square} 1/\text{lfs}(x, y) dx dy)$ . To do this, consider the following charging scheme: each square  $C_i$  is initially allocated a charge, equal to its perimeter  $|\partial C_i|$ . Then, in order from larger squares to smaller ones, we reallocate this charge by the following process: if square  $C_i$  has all four quadrants containing two or more input sites, we remove the charge from  $C_i$  and partition it equally among its four children. In this way, the total charge remains unchanged, while each square may receive charge at most equal to  $|\partial C_i|/2^j$  from its  $j$ th ancestor, so the total charge is at most  $2|\partial C_i|$ .

Next, in any square  $C_i$  with nonzero charge containing two or more sites, we reallocate the charge from  $C$  to a child  $C_j$  of  $C$  that contains zero or one sites. This child could not have been charged by its parent in the previous reallocation step, and so ends up with a total charge at most equal to  $5|\partial C_j|$ .

Finally, we observe that all charge has been concentrated in squares with zero or one sites, and that each such square has a charge proportional to its perimeter. These squares are disjoint, and within any such square, the charge is at most proportional to the integral of  $1/\text{lfs}(x, y)$ .  $\square$

LEMMA 5.2. *In any hierarchical clustering, the sum of perimeters of convex hulls of clusters is  $O(\int_{(x,y) \in \square} 1/\text{lfs}(x, y) dx dy)$ .*

*Proof.* Let  $S_i$  be the clusters in some clustering of  $S = S_0$ . For each cluster  $S_i$  in the clustering, let  $C_i$  be a minimal bounding square, let  $3C_i$  denote a concentric square three times as wide, and let  $D_i = (3C_j \setminus 3C_i) \cap \square$ , where  $S_j$  is the parent of  $S_i$  in the clustering. Further, let  $L(x, y, i) = 1/d((x, y), S_i)$  where  $d$  denotes the Euclidean distance to the closest point in cluster  $S_i$ . Then

$$\begin{aligned} \sum_i |\partial S_i| &\geq \frac{1}{2} \sum_i |\partial C_i| \\ &\geq \frac{1}{6} \sum_i |\partial 3C_i| \\ &\geq \frac{1}{12} \sum_i |\partial D_i|, \end{aligned}$$

because each bounding square has at most twice the perimeter of its cluster, and in the final sum of this sequence of inequalities, the square  $3C_i$  contributes to the boundary of two clusters  $D_j$  for its two children.

Further, for any cluster  $S_i$ ,

$$|\partial D_i| = \Omega\left(\int_{D_i} L(x, y, i) \, dx \, dy\right).$$

To show this, form a sequence of concentric squares, starting at  $C_i$  and doubling in size at each step of the sequence until the last member of the sequence contains all of  $D_i$ . Within any of the annular regions between two adjacent pairs of squares in this sequence, the integrand is inversely proportional to the perimeters of the two squares, so the integral in this same annular region is directly proportional to its outer square's perimeter. The overall integral is the sum of the integrals within each of these annular regions, a sum that is proportional to a geometric series adding to the perimeter of  $D_i$ .

Finally, observe that, for  $(x, y) \in \square$ ,

$$1/\text{lfs}(x, y) \leq L(x, y, 0) \leq \max_{(x, y) \in D_i} L(x, y, i).$$

The first equality holds because the local feature size is a distance to some point  $p$  in  $S = S_0$ , and the second holds because the clusters containing  $p$  correspond to sets  $D_i$  that together cover  $\square$ .

Putting these claims together,

$$\begin{aligned} \sum_i |\partial S_i| &= \Omega\left(\sum_i |\partial D_i|\right) \\ &= \Omega\left(\sum_i \int_{D_i} L(x, y, i) \, dx \, dy\right) \\ &= \Omega\left(\int_{\square} \sum_{(x, y) \in D_i} L(x, y, i) \, dx \, dy\right) \\ &= \Omega\left(\int_{\square} \max_{(x, y) \in D_i} L(x, y, i) \, dx \, dy\right) \\ &= \Omega\left(\int_{\square} 1/\text{lfs}(x, y) \, dx \, dy\right), \end{aligned}$$

as was to be shown.  $\square$

We note that a similar argument based on local feature sizes can be used to provide an alternative proof for our previous result [14] that quadtree triangulations have total length within a constant factor of the minimum length Steiner triangulation of a point set.

Putting these results together, we have:

**THEOREM 5.1.** *In  $O(n \log n)$  time, we can construct a hierarchical clustering with total cluster perimeter within a constant factor of optimal for any set of Euclidean points.*

### 5.3 Euclidean Pants

We now show how to convert the hierarchical clustering resulting from the method of Theorem 5.1 into a pants decomposition that approximates the optimal pants decomposition. We must be careful in doing this, as (due to the pants decomposition lacking a curve that surrounds the entire point set) the optimal pants decomposition may itself have significantly lower length than the compressed quadtree. Nevertheless, we show that the quadtree based clustering leads to a good pants decomposition.

**THEOREM 5.2.** *In  $O(n \log n)$  time, we can construct a pants decomposition with total length within a constant factor of optimal for any set of Euclidean points.*

*Proof.* Recall that Theorem 5.1 produces a clustering in which disjoint clusters have disjoint convex hulls. Let  $\epsilon$  be  $n^{-2}$  times the minimum distance between hulls of disjoint clusters in this clustering; this distance may be found in  $O(n \log n)$  time via a medial axis computation. For each cluster  $C_i$  in the clustering produced by Theorem 5.1 other than the cluster of all sites, surround  $C_i$  by a curve at distance  $|C_i|\epsilon$  from the convex hull of  $C_i$ . The total length of this curve is  $2\pi|C_i|\epsilon$  plus the length of the convex hull; these  $2\pi|C_i|\epsilon$  terms are negligible, even when added up over all the clusters, and may be ignored for the rest of our analysis. Finally, when pairing up quadrants of the root square to form the overall hierarchical clustering, do so in the way that minimizes the overall length.

We distinguish two cases in our analysis. Let  $s$  be the side length of the minimum enclosing square used as the root of our quadtree algorithm. In the first case, the optimal clustering includes a cluster with diameter at least  $s/2$ . In this case, we may add another curve surrounding the entire point set without increasing the total length of the optimal decomposition by more than a constant factor. As in Theorem 5.1, the total length of this augmented optimal decomposition is proportional to that of the decomposition found by our quadtree algorithm.

In the second case, the optimal clustering includes two subsets of small diameter, separated by a larger distance. Let  $B_1$  and  $B_2$  be the minimum bounding squares of these two subsets. Then  $B_1$  and  $B_2$  must touch opposite sides of the quadtree's root square; since they each have small diameter, no quadrant of the root square can contain points from both clusters. Therefore, all curves found by our pants decomposition are contained within (a small dilation of)  $B_1$  or  $B_2$ . We can then follow the same analysis used in the proof of Theorem 5.1 to show that the length of the quadtree-based pants decomposition is upper bounded by  $\int_{B_1 \cup B_2} 1/\text{fs}(x, y) dx dy$  while the length of the optimal pants decomposition is lower bounded by the same quantity.  $\square$

## 6 Hyperbolic Pants

We now describe analogous problems of clustering minimizing the sum of convex hull perimeters, and of optimal pants decomposition, for the hyperbolic plane instead of the Euclidean plane. Point sets in the hyperbolic plane with constant diameter can be well approximated by Euclidean point sets, so in this case we could apply our quadtree-based Euclidean approximation, but this approach does not work for more widely spaced point sets. Our eventual algorithm combines our Euclidean square pants decomposition with our tree clustering method for general metric spaces.

For simplicity of exposition we describe our algorithms as based on primitives that can compute exact distance-based predicates of points in hyperbolic space, ignoring questions of how such points are represented. However, as our eventual result is an approximation algorithm, our results can be extended

without difficulty to a computational model in which all distance computations are approximate; we omit the details.

The connection between tree clustering (in which we attempt to optimize the sum of spanning tree lengths) and the hyperbolic clustering problem (in which we instead attempt to optimize the sum of convex hull perimeters) can be made more concrete by the following lemma.

**LEMMA 6.1.** *Let  $\epsilon > 0$  be a fixed constant, and let  $S$  be a set of points in the hyperbolic plane  $\mathbb{H}^2$ , such that no two points in  $S$  are closer than distance  $\epsilon$  to each other. Then the convex hull perimeter and minimum spanning tree length of  $S$  are within constant factors of each other.*

*Proof.* Let  $T$  denote the minimum spanning tree of  $S$  (Figure 8, top left) and  $H$  denote the convex hull. In one direction, the vertices of  $H$  form a subsequence of an Euler tour of  $T$ , so (regardless of point separation)  $|H| \leq 2|T|$ .

In the other direction,  $T$  of  $S$  has length within a constant factor of the minimum Steiner tree of  $S$  (Figure 8, top right), e.g., by a similar Euler tour argument. If we let  $S'$  be a maximal set of points having the same convex hull as  $S$  and having no pair closer than  $\epsilon$ , and  $T'$  be a minimum spanning tree of  $S'$ , then the length of  $T'$  is again at least a constant factor times that of  $T$ , as adding points can not decrease the Steiner tree length (Figure 8, center left) and the minimum spanning tree of  $T'$  has length at least that of its Steiner tree (Figure 8, center right). Note that all edges of  $T'$  have length at least  $\epsilon$  and at most  $2\epsilon$ , as if there were any longer edge we could add its midpoint to  $S'$  contradicting the assumption that  $S'$  is maximal. Thus, the total length of  $T'$  (and therefore also that of  $T$ ) is  $O(|S'|)$ .

Now, form a disk of radius  $\epsilon/2$  around each point of  $S'$ . These disks are disjoint, and all lie within the set  $C$  consisting of the points within distance  $\epsilon/2$  of the convex hull of  $S$  (Figure 8, bottom left). Therefore, the area of  $C$  is  $\Omega(|S'|)$ , and by a hyperbolic isoperimetric theorem [28], the perimeter of  $C$  is  $\Omega(|S'|)$ .

Finally, we bound the perimeter of  $C$ . In Euclidean geometry, the perimeter of a set  $C$  formed by expanding a hull  $H$  by distance  $\epsilon/2$  can be calculated exactly, as the perimeter of  $H$  plus that of a ball of radius  $\epsilon/2$  (additivity of perimeters for Minkowski sums). In hyperbolic geometry, things are more complicated, but we may still relate the perimeters of  $C$  and  $H$  as follows. Partition the perimeter of  $H$  into curves of length  $O(1)$ , and assign each point of the perimeter of  $C$  to the nearest such curve of  $H$  (Figure 8, bottom right). Each curve is assigned only to points within distance  $\epsilon/2$  of it, so the total assigned length to each curve is  $O(1)$ . Therefore,  $|H| = \Omega(|C|)$ . Putting these steps together,  $|T| = O(|T'|) = O(|S'|) = O(|C|) = O(|H|)$ .  $\square$

Thus, if all pairwise distances in our point set are  $\Omega(1)$ , our tree-based approximation for general metric spaces provides also a constant factor approximation to minimizing the sum of cluster hull perimeters in the hyperbolic plane. On the other hand, if all pairwise distances are  $O(1)$ , the Klein model of the hyperbolic plane (in which hyperbolic points are modeled as points in a Euclidean disk, and hyperbolic lines are modeled by straight chords of the disk) provides a convexity-preserving Euclidean approximation of the point set, and applying our quadtree-based approximation to this model leads to a constant factor approximation to the minimum sum of cluster hull perimeters in the hyperbolic plane. A more challenging situation arises when we are asked to cluster a point set that combines both large and small distances; we show below how to combine these two clustering approaches in an approximation algorithm for this more general case.

Our hyperbolic clustering algorithm, in rough outline, follows the following steps.

1. Find a subset  $S'$  of the input sites, such that  $S'$  has no two sites within some constant distance bound of each other, and maximal with respect to this constraint. Group the remaining sites into clusters according to the nearest member of  $S'$  to each site.




Figure 8: Steps in the proof of Lemma 6.1. The minimum spanning tree  $T$  of  $S$  (top left) is proportional in length to its minimum Steiner tree (top right). Augmenting  $S$  to a maximal  $\epsilon$ -separated set  $S'$  does not decrease the Steiner tree length (center left), which in turn is less than the length of the minimum spanning tree  $T'$  of  $S'$  (center right). The  $|S'| - 1$  edges of  $T'$ , each of length  $\Theta(1)$ , have a total length proportional to the total area of a collection of radius- $\epsilon/2$  disjoint disks around each point of  $S'$  (bottom left), which in turn is bounded by the area of a set  $C$  formed by expanding the hull of  $S$  by distance  $\epsilon/2$ . By hyperbolic isoperimetry, the area of  $C$  is bounded by its perimeter, and by partitioning the perimeter of the convex hull  $H$  of  $S$  into curves of length  $\Theta(1)$  and assigning each point of  $C$  to the nearest curve we may show that the perimeter of  $C$  is proportional to that of  $H$  (bottom right). The trees depicted are for illustrative purposes only and may not be the actual optima, as in any case the objects in the figure should be interpreted as belonging to the hyperbolic plane rather than the Euclidean plane.




Figure 9: Squarepants in a tree: a schematic view of our hyperbolic clustering. We group sites into bounded-diameter subsets, apply our Euclidean quadtree pants decomposition method to each subset, and connect the subsets with clusters following minimum spanning tree edges from our general metric space clustering algorithm.

2. Within each cluster, approximate the sites by a point set in the Euclidean plane, and apply our quadtree-based algorithm to find an approximately-optimal clustering for these points.
3. Find the minimum spanning tree of  $S'$ , and use our tree-based clustering algorithm for general metric spaces to find a clustering for  $S'$ . Adjust the boundaries of this clustering so that they surround the smaller clusters associated with each site of  $S'$ .

Figure 9 gives a schematic view of the clustering produced by this algorithm: a collection of quadtrees, connected by clusters that follow minimum spanning tree edges.

### 6.1 Finding a Well-Separated Subset

We begin the detailed description of our hyperbolic clustering algorithm by showing how to implement efficiently its first step, in which we find a well-separated subset of the input sites. It is tempting to apply existing Euclidean methods for similar problems [18], but these appear to be based on recursive subdivision of rectangles, a concept that makes little sense hyperbolically.

**LEMMA 6.2.** *If we are given as input a set  $S$  of  $n$  point sites in  $\mathbb{H}^2$ , and a constant  $\delta$ , we can find in time  $O(n \log n)$  a subset  $S' \subset S$  such that the closest pair of points in  $S'$  are at least at distance  $\delta$  apart from each other, and such that every point of  $S$  is within distance  $\delta$  of some point in  $S'$ . In the same time bound we may also find the nearest point in  $S'$  to each point in  $S$ , and list all pairs of points of  $S'$  with distance less than  $2\delta$  of each other.*






Figure 10: Illustration of the algorithm in Lemma 6.2 for finding a well-separated subset of the input sites. For each site  $p$  that we consider (for instance the site marked by a dark disk) we need only compare its distances to the already-chosen sites (hollow disks) within a bounded radius region  $R_{i,p}$  in the same annulus and  $R_{i-1,p}$  in the next inner annulus (both shaded).




Figure 11: Partitioning a set of sites into convex subsets and clustering each subset can have significantly greater total length than a clustering of the overall set.

*Proof.* We choose an arbitrary point as the origin of our hyperbolic plane, sort the sites by their distance to the origin, and group the sites into subsets, where  $S_i$  consists of the sites with distance between  $i\delta$  and  $(i+1)\delta$  from the origin. Let  $A_i$  denote the annulus containing  $S_i$ , bounded by circles with radii  $i\delta$  and  $(i+1)\delta$ ; the circles in Figure 10 depict the boundaries of these annuli.

Within each annulus  $A_i$ , we sort the sites in  $S_i$  again, in clockwise order by the angle they form with the origin. We will consider the points in this order, adding each site to  $S'$  exactly when it is at least  $\delta$  in distance from all previously-added sites. In this way, we will generate a maximal subset of sites at distance  $\delta$  or more from each other, as desired.

When we consider site  $p$ , define region  $R_{i,p}$  as the smallest annular wedge containing the intersection of  $A_i$  with a radius- $\delta$  disk centered at  $p$ ; then only sites within  $R_{i,p} \cup R_{i-1,p}$  (the shaded region in Figure 10) may be within distance  $\delta$  of  $p$ . This region has radius at most  $2\delta$ , as each point in it may be connected by a radial segment of length  $\delta$  to the radius- $\delta$  disk; since  $\delta$  is assumed to be bounded, and the sites in  $S'$  are a bounded distance apart,  $|(R_{i,p} \cup R_{i-1,p}) \cap S'| = O(1)$ . That is, by scanning sequentially a constant number of steps from  $p$  forwards and backwards in the sorted orders for  $S' \cap A_i$  and  $S' \cap A_{i-1}$ , until each sequential search reaches a point outside  $R_{i,p} \cup R_{i-1,p}$ , we may find any site in  $S'$  that can be within distance  $\delta$  of  $p$ . If one of these  $O(1)$  candidate neighbors has distance less than  $\delta$ , we eliminate  $p$ ; otherwise, we add  $p$  to  $S'$ .

The time for this algorithm is  $O(n \log n)$  for the two sorting steps, and then  $O(1)$  per point to march stepwise through the sorted orders for  $S_i$  and  $S' \cap A_{i-1}$  and determine whether to keep or eliminate the point, for a total time of  $O(n \log n)$ .

Once we have constructed  $S'$ , finding the nearest neighbor in  $S'$  for each point in  $S$  may be done by a similar scan through a constant number of sites in  $(R_{i,p} \cup R_{i-1,p} \cup R_{i+1,p}) \cap S'$ , and listing pairs of sites in  $S'$  within distance  $2\delta$  of each other may be performed via a similar search of five neighboring annuli for each site.  $\square$

We leave as an open problem efficiently solving the generalized version of this problem in which  $\delta$  is not restricted to be a constant; the difficulty with this generalization is that the regions  $R_{i,p}$  as defined in the proof above may contain nonconstant numbers of sites of  $S'$ .

## 6.2 Clustering Low-Diameter Neighborhoods

Form the Voronoi diagram of the sites in the subset  $S'$  described in Lemma 6.2. For any site  $c_i \in S'$ , let  $V_i$  be formed as the intersection of the Voronoi cell of  $c_i$  with a disk of radius  $\delta$  centered on  $c_i$ . Each cell  $V_i$  is a convex and has bounded diameter. Therefore, an optimal clustering for the sites in  $V_i$  may be approximated easily enough, as follows.

Embed  $V_i$  in a Klein model of the hyperbolic plane; that is, a unit disk in the Euclidean plane, in which hyperbolic lines are modeled as straight line segments through the disk. This model has the convenient property that convex hulls in the hyperbolic plane are modeled by convex hulls of the corresponding points in the model. Choose the embedding in such a way that  $c_i$  is placed at the center of the disk. As the sites in  $V_i$  have bounded distance from the center, the distortion of the embedding is also bounded. Then, use our quadtree based algorithm to approximate the optimal clustering of the embedded sites. By the bound on the distortion, the result is also an approximation to the optimal clustering of the same sites in the original hyperbolic plane.

However, we need a stronger conclusion. It is not enough that each cell  $V_i$  is clustered approximately optimally. Rather, we need the sum of the cluster lengths in all cells to approximate the optimal clustering of the entire set. To see that this is a nontrivial requirement, see Figure 11: there exists a point set, and a partition of the point set into convex regions, such that the sum of the optimal clustering lengths within each region is much larger than the optimal clustering of the overall point set. However, as we now show, such a phenomenon can not occur with our partition into Voronoi cells due to the low aspect ratio of these cells.

**LEMMA 6.3.** *Let  $C$  be a simple closed curve, and  $T$  be the sites inside  $C$ . Then there exists a collection of curves  $C_i$ , each consisting of a disjoint union of zero or more simple closed curves, such that each  $C_i$  encloses  $T \cap V_i$ , and such that the total length of all the curves  $C_i$  is proportional to a constant times the length of  $C$ .*

*Proof.* We construct the curves  $C_i$  by intersecting  $C$  with each cell  $V_i$ , and reconnecting each end of the resulting curves by a curve passing around the boundary of cell  $V_i$ , as shown in Figure 12.

We must show that the additional length created in this reconnection process is at most proportional to the length of  $C$ . We distinguish the components of the resulting collection of curves into two cases: components containing the center  $c_i$  of their cell  $V_i$ , and all other components.

If  $C$  contains  $k$  centers of cells, and passes through two or more cells, then by Lemma 6.1 the length of  $C$  must be  $\Omega(k)$ . Within each cell  $V_i$ , the total length of the connection between the endpoints of the component of  $C$  containing  $c_i$  is  $O(1)$ , so the total length in all such cells is  $O(k)$ . Therefore, the total reconnection length for these components is proportional at most to that of  $C$ .




Figure 12: Any curve passing through the cells  $V_i$  can be divided into components in each cell, increasing the total length by at most a constant factor (Lemma 6.3).




Figure 13: Worst-case configuration for Lemma 6.3: a path of length  $\delta$  connects two nearby points on the boundary, avoiding  $c_i$ , so when forming  $C_i$  we must use nearly all of the boundary of cell  $V_i$ .

Finally, suppose  $C_i$  is a component of the collection of curves, not containing  $c_i$ , and let  $r$  be the ratio between the length of the portion of the boundary of  $V_i$  bounding  $C_i$  and the portion of  $C$  bounding  $C_i$ . The maximum possible value of  $r$  occurs when  $C$  connects the two points where  $C$  crosses the boundary of  $V_i$  by as short a path as possible (avoiding  $c_i$ ), and where the portion of boundary between these two crossing points is as long as possible: that is, except possibly for the straight line segments containing the two crossing points, the boundary of  $V_i$  follows the radius- $\delta$  disk centered around  $c_i$ . Next, to maximize  $r$ , the length of the path formed by intersecting  $C$  with  $V_i$  should be as short as possible. That is, the two endpoints must be at distance exactly  $\delta/2$  from  $c_i$ , as no boundary point of  $V_i$  may be closer to  $c_i$  than that. Finally, among all ways of choosing two endpoints at distance  $\delta/2$  from  $c_i$ , the one maximizing  $r$  is the one where the two endpoints are coincident, and the portion of  $C_i$  following the boundary of  $V_i$  follows the entire boundary (Figure 13). For this worst-case configuration,  $r$  is constant whenever  $\delta$  is constant. Therefore, the total reconnection length for components of this section type is again proportional at most to that of  $C$ .  $\square$

**COROLLARY 6.1.** *If we are given as input a set of sites in  $\mathbb{H}^2$ , group them into cells  $V_i$  as described above, use the Klein model to map each cell with bounded distortion in the Euclidean plane, and apply our quadtree clustering method to the resulting family of sets of Euclidean sites, then the total length of the resulting clusterings is at most proportional to the total length of the optimal clustering of all the sites.*

*Proof.* From the optimal clustering of all the points we can form a clustering within each  $V_i$  by restricting

the clusters to the sites within  $V_i$ . By applying Lemma 6.3 to the convex hulls of each cluster, we see that the total length of the convex hulls of each of these clusterings is at most proportional to the total length of the original optimal clustering. Since the quadtree clustering approximates the optimal clustering within each cell  $V_i$ , the sum of the quadtree clusterings is again within a constant factor of the overall optimal clustering.  $\square$

### 6.3 Clusters of Clusters

We now show how to find larger clusters in our clustering, connecting multiple sites of  $S'$ . Our algorithm for this part follows that for general metric spaces: we form the minimum spanning tree of  $S'$ , expand it by zero-length edges so that all internal vertices have degree three, and repeatedly split the resulting tree on the edge that most evenly balances the weights of the two subtrees formed by the split.

**LEMMA 6.4.** *Let  $S$ ,  $S'$ , and the cells  $V_i$  be as described above. Then in  $O(n \log n)$  time, we can find a noncrossing family of curves, including among them the convex hulls of the sites in each cell  $V_i$ , such that each component of the plane minus these curves is either one of these convex hulls or a pair of pants. The total length of these curves is proportional to that of the optimal clustering of  $S$ .*

*Proof.* We note that the optimal clustering of  $S$  has length at least that of the optimal clustering of  $S'$ , as we can omit redundant boundary curves from any clustering of  $S$  to form a clustering of  $S'$ . Further, by Lemma 6.1, the length of the optimal clustering of  $S'$  in terms of the sum of convex hulls of clusters is proportional to the length of the optimal clustering in terms of the sum of spanning tree lengths of clusters.

As in our general metric clustering algorithm, we form the minimum spanning tree of  $S'$ , expand it by zero-length edges so that all internal vertices have degree three, and repeatedly split the resulting tree on the edge that most evenly balances the weights of the two subtrees formed by the split. In order to reduce crossings between clusters for our application of this method to pants decomposition, whenever we split a vertex, we choose a split consistent with the radial ordering of the neighbors to that vertex, so that the expanded tree could be viewed as embedded without crossings in the hyperbolic plane, with the points into which each original vertex has been expanded placed within a small disk near the original point location.

We then cluster this expanded tree by greedy splitting as in the general metric clustering algorithm. We surround each cluster of the clustering by a curve that follows the boundary of the hull of a cell  $V_i$  whenever the subtree corresponding to that cluster passes through  $c_i$ , and that (outside these hulls) follows the Euler tour of the minimum spanning tree of  $S'$ .

If the subtree corresponding to the cluster passes through  $k$  hulls, its length is increased by  $O(k)$ , but (by Lemma 6.1) the length of the convex hull of the cluster must already have been  $\Omega(k)$ . Therefore, the total length of this system of curves is proportional to that of the clustering of  $S'$ , which (by Theorem 4.1) is proportional to that of the optimal clustering.  $\square$

### 6.4 Hyperbolic Clustering and Pants

We are now ready to put these results together, in our overall hyperbolic clustering algorithm.

**THEOREM 6.1.** *In  $O(n \log n)$  time, we can construct a hierarchical clustering with total cluster perimeter within a constant factor of optimal for any set of hyperbolic sites.*

*Proof.* We construct a set of  $\delta$ -separated points by Lemma 6.2, form the Voronoi cells  $V_i$  (each such cell being formed as the intersection of a disk with a constant number of halfplanes, the neighbors of each center  $c_i$  being found as described in Lemma 6.2), and assign each site to the nearest center  $c_i$  as

described in Lemma 6.2. We then form clusters within each cell  $V_i$  as described in Corollary 6.1 and clusters grouping sets of cells as described in Lemma 6.4.  $\square$

The same methods and proof apply, essentially unchanged, for hyperbolic pants decomposition. The quadtree clustering used within each cell  $V_i$  can be modified as in the Euclidean case to produce a family of noncrossing curves with essentially the same length, and Lemma 6.4 already describes how to form noncrossing curves for its clusters.

**THEOREM 6.2.** *In  $O(n \log n)$  time, we can construct a pants decomposition with total length within a constant factor of optimal for any set of hyperbolic sites.*

## 7 Conclusion

We have provided efficient approximation algorithms for clustering in general metric spaces and the Euclidean and hyperbolic planes. We suspect that our NP-completeness proof for general metric clustering can be extended to show that it is hard to approximate the problem within a factor better than  $1 + \epsilon$ , for some fixed  $\epsilon$ , but it would nevertheless be of interest to reduce the 3.42 approximation ratio of our algorithm.

For the Euclidean and hyperbolic problems, we do not know whether finding the optimal clustering is intractable. Also, the approximation ratios of our algorithms for these problems are large and inexplicit. It would be of interest to resolve the computational complexity of these problems, and to find improved approximations for them.

Finally, there are many other objective functions for clustering quality than the ones we have considered. As an example, it is natural to consider clustering planar point sets by convex hull area instead of perimeter; this version of the problem is closely related to covering point sets by few lines [1,2], as collinear point sets lead to zero-area clusters. Much work remains to be done on algorithms for optimizing or approximating optimal clusterings for this and other criteria.

## Acknowledgements

A preliminary version of this paper appeared in the Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms, 2007. We thank Mohammadreza Ghodsi for pointing out an error in an earlier version of the proof of Lemma 5.2.

## References

- [1] P. K. Agarwal and C. M. Procopiuc. Approximation algorithms for projective clustering. *J. Algorithms* 46(2):115–139, 2003.
- [2] P. K. Agarwal, C. M. Procopiuc, and K. R. Varadarajan. A  $(1 + \epsilon)$ -approximation algorithm for 2-line-center. *Computational Geometry Theory and Applications* 26(2):119–129, 2003.
- [3] S. Alstrup, J. Holm, and M. Thorup. Maintaining center and median in dynamic trees. *Proc. 7th Scand. Worksh. Algorithm Theory (SWAT 2000)*, pp. 46–56. Springer-Verlag, Lecture Notes in Computer Science 1851, 2000.
- [4] S. Aluru and F. E. Sevilgen. Dynamic compressed hyperoctrees with application to the  $N$ -body problem. *Proc. 19th Conf. Found. Softw. Tech. Theoret. Comput. Sci.*, pp. 21–33. Springer-Verlag, Lecture Notes in Computer Science 1738, 1999.
- [5] N. Amenta, M. W. Bern, and D. Eppstein. The crust and the  $\beta$ -skeleton: combinatorial curve reconstruction. *Graphical Models & Image Processing* 60/2(2):125–135, March 1998.
- [6] J. L. Bentley, D. D. Sleater, R. E. Tarjan, and V. K. Wei. A locally adaptive data compression scheme. *Commun. ACM* 29(4):320–330, 1986.
- [7] M. W. Bern. Approximate closest-point queries in high dimensions. *Inform. Process. Lett.* 45:95–99, 1993.

- [8] M. W. Bern, D. Eppstein, and S.-H. Teng. Parallel construction of quadrees and quality triangulations. *Proc. 3rd Worksh. Algorithms and Data Structures (WADS 1993)*, pp. 188–199. Springer-Verlag, Lecture Notes in Computer Science 709, 1993.
- [9] C. Blundo and R. de Prisco. New bounds on the expected length of one-to-one codes. *IEEE Trans. Info. Theory* IT-42:246–250, 1996.
- [10] A. Borodin, R. Ostrovsky, and Y. Rabani. Subquadratic approximation algorithms for clustering problems in high dimensional spaces. *Machine Learning* 56:153–167, 2004.
- [11] K. L. Clarkson. Fast algorithms for the all nearest neighbors problem. *Proc. 24th IEEE Symp. Foundations of Computer Science (FOCS 1983)*, pp. 226–232, 1983.
- [12] S. Dasgupta. Performance guarantees for hierarchical clustering. *Proc. 15th Conf. Computational Learning Theory (COLT 2002)*, pp. 351–363. Springer-Verlag, Lecture Notes in Computer Science 2375, 2002.
- [13] P. Elias. Interval and recency rank source coding: two on-line adaptive variable-length schemes. *IEEE Trans. Info. Theory* IT-33:3–10, 1987.
- [14] D. Eppstein. Approximating the minimum weight Steiner triangulation. *Discrete & Computational Geometry* 11(2):163–191, 1994.
- [15] D. Eppstein. Fast hierarchical clustering and other applications of dynamic closest pairs. *J. Experimental Algorithmics* 5(1):1–23, June 2000, <http://www.jea.acm.org/2000/EppsteinDynamic/>, arXiv:cs.DS/9912014.
- [16] D. Eppstein, M. T. Goodrich, and J. Z. Sun. The skip quadtree: a simple dynamic data structure for multidimensional data. *Proc. 21st Symp. Computational Geometry*, pp. 296–305. ACM, 2005, arXiv:cs.CG/0507049.
- [17] J. Erickson and É. Colin de Verdière. Tightening non-simple paths and cycles on surfaces. *Proc. 17th ACM-SIAM Symp. Discrete Algorithms (SODA 2006)*, pp. 192–201, 2006.
- [18] T. Feder and D. H. Greene. Optimal algorithms for approximate clustering. *Proc. 20th ACM Symp. Theory of Computing (STOC 1988)*, pp. 434–444, 1988.
- [19] M. R. Garey and D. S. Johnson. *Computers and Intractability: a Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [20] S. C. Johnson. Hierarchical clustering schemes. *Psychometrika* 32(3):241–254, 1967.
- [21] D. G. Kirkpatrick and P. Hell. On the complexity of a generalized matching problem. *Proc. 10th ACM Symp. Theory of Computing (STOC 1978)*, pp. 240–245, 1978.
- [22] R. Krauthgamer and J. R. Lee. Algorithms on negatively curved spaces. *Proc. 47th IEEE Symp. Foundations of Computer Science*, 2006.
- [23] S. K. Leung-Yan-Cheong and T. M. Cover. Some equivalences between Shannon entropy and Kolmogorov complexity. *IEEE Trans. Info. Theory* IT-24:331–338, 1978.
- [24] S.-H. Poon and S. Thite. Pants decomposition of the punctured plane. ACM Computing Research Repository, arXiv:cs.CG/0602080.
- [25] J. Ruppert. A Delaunay refinement algorithm for 2-dimensional mesh generation. *J. Algorithms* 18(3):548–585, May 1995.
- [26] N. Saitou and M. Nei. The neighbor-joining method: a new method for reconstructing phylogenetic trees. *Molecular Biology and Evolution* 4(4):406–425, 1987.
- [27] C. E. Shannon. A mathematical theory of communication. *Bell System Technical J.* 27:379–423, 1948.
- [28] E. Teufel. A generalization of the isoperimetric inequality in the hyperbolic plane. *Archiv der Mathematik* 57(5):508–513, 1991.
- [29] A. D. Wyner. An upper bound on the entropy series. *Inform. Contr.* 20:176–181, 1972.
- [30] J. Zupan. *Clustering of Large Data Sets*. Chemometrics Research Studies Ser. Research Studies Press, Chichester, 1982.